Bank of Canada Banque du Canada

Working Paper 2006-46 / Document de travail 2006-46

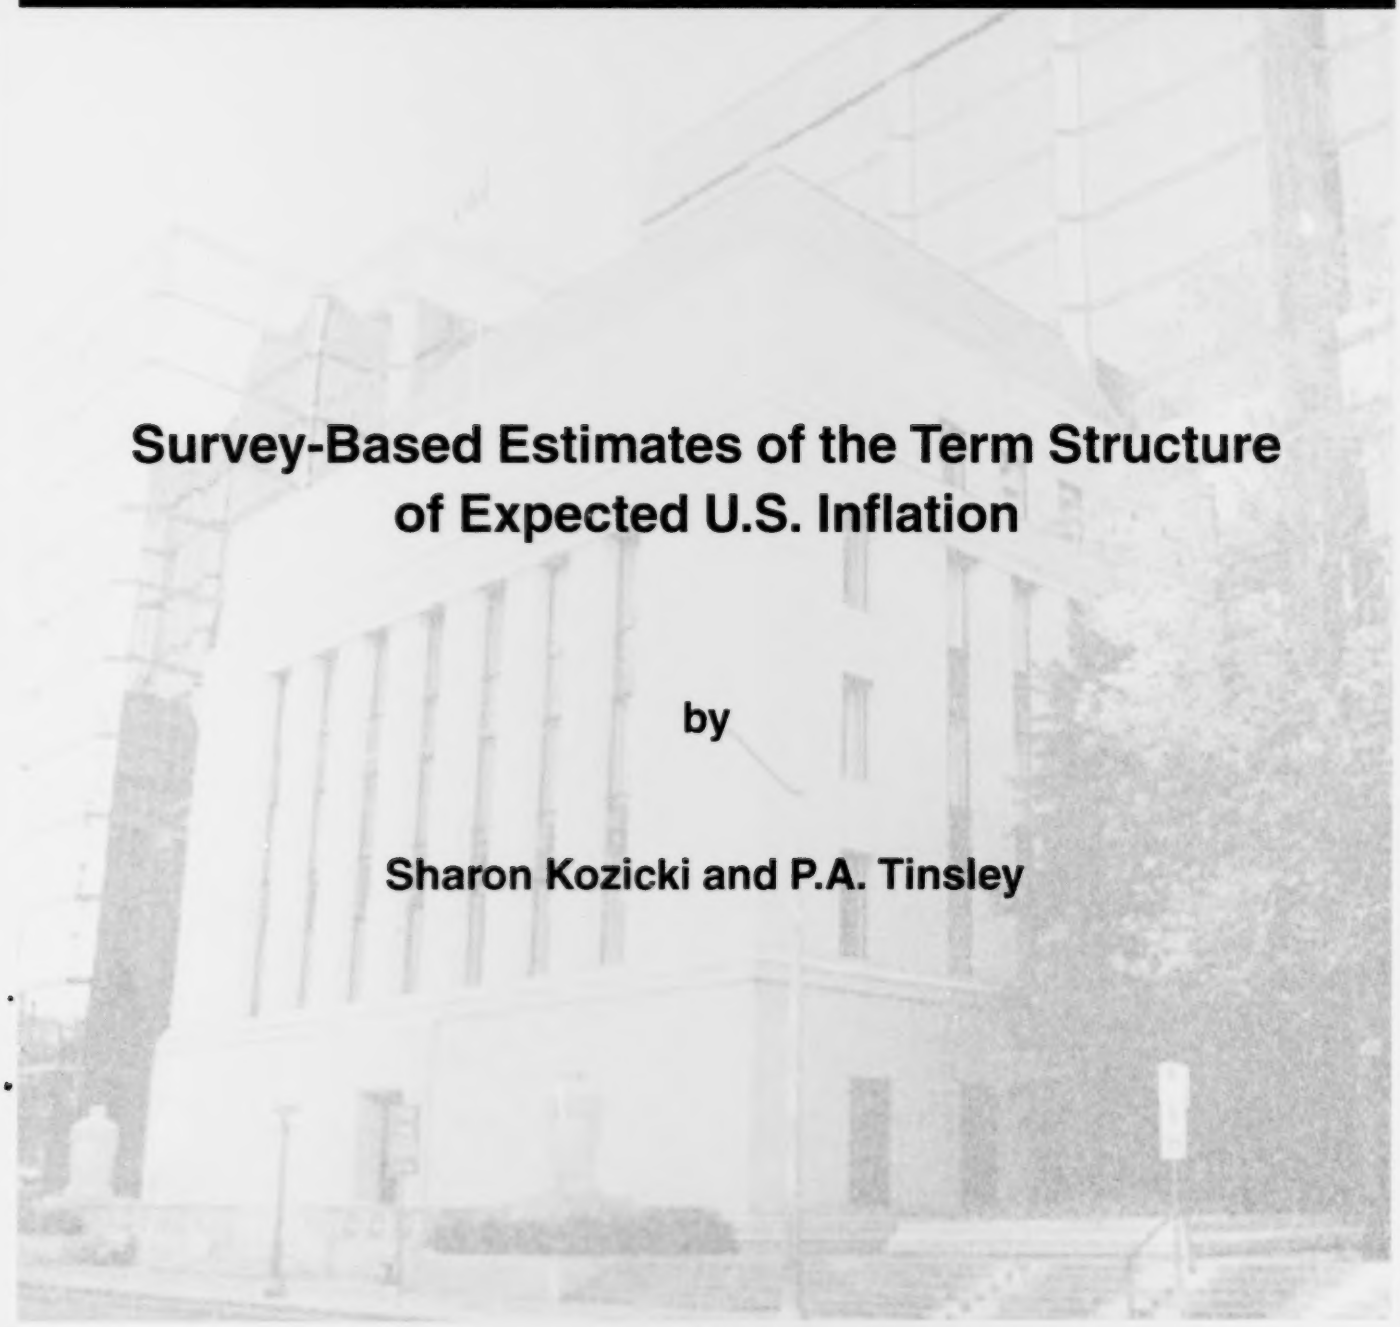

Survey-Based Estimates of the Term Structure of Expected U.S. Inflation

by

Sharon Kozicki and P.A. Tinsley

ISSN 1192-5434

Printed in Canada on recycled paper

2008-02027/

Bank of Canada Working Paper 2006-46

December 2006

Survey-Based Estimates of the Term Structure of Expected U.S. Inflation

by

Sharon Kozicki[1] and P.A.Tinsley[2]

[1]Research Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
skozicki@bankofcanada.ca

[2]Birkbeck College
University of London

The views expressed in this paper are those of the authors.
No responsibility for them should be attributed to the Bank of Canada.

Contents

Acknowledgements iv
Abstract/Résumé v

1. Introduction 1

2. A Model of Survey Expectations 3
   2.1 A shifting-endpoint AR model 4
   2.2 Approximating survey expectations with AR expectations 7
   2.3 A state space model of the inflation endpoint 8
   2.4 Dealing with missing observations 10

3. Empirical Results 11
   3.1 Data 11
   3.2 Results 12

4. Heterogeneous Perceptions of Inflation Targets 17

5. Concluding Comments 18

References 20

Tables 23

Figures 27

Acknowledgements

We are grateful for comments received from seminar participants at the University of Kansas, the Federal Reserve Bank of Kansas City, Carleton University, and attendees of the 2006 AEA session on "Uncertainty and Monetary Policy and the 2006 Canadian Economic Association session on Inflation Dynamics. Special thanks to Jean Boivin, Todd Clark, Andrea Raffo, and Francisco Ruge-Murcia for insights that have considerably improved the article. We are indebted to Matthew Cardillo for excellent research assistance.

Abstract

Surveys provide direct information on expectations, but only short histories are available at quarterly frequencies or for long-horizon expectations. Longer histories typically contain only semi-annual observations of short-horizon forecasts. The authors fill in the gaps by constructing a 50-year monthly history of expected inflation at all horizons from one month to 10 years that is consistent with inflation data and infrequent survey data. In the process, some models that fit inflation well are found to generate forecasts that bear little resemblance to survey data. Also, survey data on near-term expectations are found to contain considerable information about long-horizon views. The estimated long-horizon forecast series, a measure of the private sector's perception of the inflation target of monetary policy, has shifted considerably over time and is the source of some of the persistence of inflation. When compared with estimates of the effective inflation goal of policy, these perceptions suggest that monetary policy has been less than fully credible historically.

JEL classification: E3, E5
Bank classification: Inflation and prices; Inflation targets; Uncertainty and monetary policy

Résumé

Les enquêtes offrent une information de première main sur les attentes d'inflation, mais celles de périodicité trimestrielle ou portant sur des horizons éloignés sont encore jeunes. Sur de plus longues périodes, on ne dispose en général que d'observations recueillies à une fréquence semestrielle et se rapportant à un horizon rapproché. Pour suppléer aux observations manquantes, les auteurs élaborent une série mensuelle qui mesure sur 50 ans les attentes d'inflation à tous les horizons compris entre un mois et dix ans et qui est conforme aux données de l'inflation et aux résultats des enquêtes menées à intervalles peu fréquents. Ce faisant, ils constatent que certains modèles qui décrivent pourtant bien l'évolution passée de l'inflation génèrent des prévisions peu conformes aux données d'enquête. Ils remarquent également que les données d'enquête sur les attentes à court terme renferment une masse considérable d'informations sur les attentes à long terme. Leur estimation des taux d'inflation prévus aux horizons éloignés, qui donne une idée des perceptions des agents du secteur privé au sujet de la cible d'inflation des autorités monétaires, a beaucoup changé au fil du temps et explique en partie la persistance de l'inflation. Confrontées aux estimations de l'objectif d'inflation réel des autorités, ces perceptions donnent à penser que la politique monétaire n'était pas parfaitement crédible dans le passé.

Classification JEL : E3, E5
Classification de la Banque : Inflation et prix; Cibles en matière d'inflation; Incertitude et politique monétaire

1. Introduction

Information on expected inflation at short and long horizons is key to assessing the credibility of monetary policy, to examining how borrowing decisions of households and firms respond to shifts in real costs of debt, and to evaluating the expected inflation response to monetary policy actions. Unfortunately, direct observations on market expectations of inflation are limited.

Surveys of forecasts provide one source of direct information on expectations.[1] However, surveys are infrequently used, partly due to the incomplete sampling design of available surveys: only short time series are available for surveys that sample at quarterly frequencies or higher. In addition, lengthy time series are available only for surveys of short-horizon forecasts, generally two- or four-quarter outlooks, and are often collected only at semi-annual intervals.

These limitations frequently lead researchers to alternative proxies for inflation expectations.[2] Some analyses use forecasts from econometric models.[3] Others extract estimates of average expected inflation from interest rate data.[4] However, these proxies may not resemble the expectations revealed in surveys, calling into question inferences drawn from these proxies regarding policy credibility, investment decision-making, and monetary policy transmission.

This paper addresses the limitations to survey data and other inflation-expectations proxies by constructing a 50-year history of monthly ex ante measures of expected inflation and a *term structure of expected inflation* for the United States. The constructed measures of

[1]Historically, based on the argument that survey participants have no incentives to provide their true expectations, some analysts have argued that surveys may not be good measures. However, the superior forecasting performance of surveys documented by Ang, Bekaert, and Wei (2005) casts some doubt on such concerns.

[2]Examples of studies that have directly used survey data to measure expected inflation include Roberts (1995 and 1997) and Kozicki and Tinsley (2002).

[3]Harvey (1988) forecasts inflation using an IMA(1,1) model to construct an expected inflation series. Laubach and Williams (2003) proxy inflation expectations with the forecast of the four-quarters-ahead percentage change in the price index for personal consumption expenditures, excluding food and energy generated from an AR(3) of inflation estimated over the prior 40 quarters.

[4]Expected inflation measures can also be constructed using nominal and indexed bonds (Breedon and Chadha 1997; Söderlind and Svensson 1997), but require assumptions on the term premium and relative liquidity of the assets. Shen and Corning (2001) and Côté et al. (1996) discuss, respectively for U.S. and Canadian data, distortions in measures of inflation constructed as the difference between yields on nominal and real yields. Moreover, using Canadian data, Christensen, Dion, and Reid (2004) find that the break-even inflation rate (BEIR), defined as the difference between nominal and real return bond yields, is, on average, higher and more variable than survey measures of expected inflation, and they argue that the risk premium and other distortions account for these observations. Consequently, they conclude that the BEIR is not a good gauge of the credibility of monetary policy.

expected inflation (time t forecasts of inflation in $t+h$) are for short and long horizons, and fill in "holes" in observations (t) and horizons (h). In addition, the constructed measures provide good fits of available survey observations.

The paper uses inflation data and Livingston Survey data on inflation expectations. The survey data ensure consistency of the constructed forecasts with such measures of expectations. The estimation uses a time-varying forecast methodology that assumes the unobserved cross-section of expectations formulated in a given period is consistent with recent inflation and available survey data on expectations. Estimates of the time-varying term structures of inflation appear to be relatively robust to the pattern of missing observations in historical survey data. That is, term structures constructed only on the basis of inflation and short-horizon survey expectations are close to those that also use longer-horizon survey expectations. Moreover, long-horizon constructions are close to long-horizon survey data, even when the latter are not used during estimation, suggesting that relatively short-horizon forecasts provide considerable information on long-horizon views.

One use of the constructed forecasts of long-horizon inflation is to examine the historical credibility of monetary policy. Inflation expectations are generally anchored by private sector perceptions of the central bank's inflation target. The article examines the historical credibility of monetary policy by comparing private sector perceptions with estimates of the "effective" inflation target of U.S. monetary policy.

This paper is organized as follows. Section 2 describes the empirical model used to approximate survey expectations. The same shifting-endpoint model is assumed to fit inflation and generate the survey data. Expressions for inflation and survey expectations are set in a state space framework, so that the unobserved perceived inflation target that anchors long-horizon expectations can be estimated. The methodology is adapted to accommodate different observation frequencies of inflation (monthly) and survey data (semi-annually), as well as missing observations of long-horizon expectations for most of the survey sample. Section 3 reports empirical results. A monthly term structure of expected inflation is constructed using the estimated model. Estimates of long-horizon expectations are consistent with constructions based on other data sets and different methodologies, as well as with available survey data (including both survey data used during estimation and survey data from other sources not used during estimation). In section 4, a comparison of the perceived inflation target with estimates of the central bank's effective inflation target provides strong evidence of heterogeneous expectations. Large differences in the 1980s suggest less than full credibility of low-inflation policy objectives. More recent convergence signals an improvement in

credibility. Section 5 offers some conclusions.

2. A Model of Survey Expectations

Survey data on expectations provide extra information that is often ignored by empirical researchers. However, survey participants implicitly provide information on their beliefs about how the economy operates. While some participants may report forecasts generated by unadjusted econometric models of the U.S. economy, most incorporate judgment into their views about what they expect the future to bring.[5] Such forecasts tend to reflect information that is not well summarized by historical data or econometric equations. Examples include structural changes, such as changes in tax laws, perceived shifts in the long-run inflation goals of policy, or changes in perceptions of policy credibility. One important characteristic of perceived structural change is that, just as it can immediately be incorporated into judgment, it will tend to immediately influence forecasts, including long-horizon forecasts.

Such perceived structural changes are often not well captured in standard empirical proxies for expectations. Reduced-form time-series models such as vector autoregressions (VARs) are popular specifications that are easy to use in multi-period forecasting exercises, owing to their linearity. They do not require practitioners to take a stand on the underlying structural model, yet perform relatively well over short horizons.[6] However, their ability to effectively accommodate structural change is limited. For instance, one approach to introducing the prospect for structural change into VAR models is to allow all model coefficients to change.[7] However, this approach tends to lead to in-sample overfitting problems and poor out-of-sample forecasting performance.

For the application in this paper, the main difficulty encountered with both univariate and multivariate autoregressive specifications is that they tend to generate multi-period forecasts that do not resemble available survey data on expectations (Kozicki and Tinsley 1998, 2001a, b). In particular, long-horizon forecasts of inflation from mean-reverting specifications are

[5] Wallis (1989) surveys developments in macroeconomic forecasting, including a discussion of judgmental forecasts as well as structural and time-series models. Sims (2002) discusses forecasting exercises at several central banks, and offers commentary on the role of "'subjective' forecasting based on data analysis by sectoral 'experts'." See also Reifschneider, Stockton, and Wilcox (1997) for the use of judgment with econometric models in the Federal Reserve's monetary policy process.

[6] McNees (1986) provides evidence that forecasts from Bayesian VARs are among the most accurate for forecasting several key U.S. macroeconomic variables. That said, Wallis et al. (1986, 1987) find that, for U.K. data, VAR forecasts do not dominate model-based forecasts.

[7] A simple approach taken by some researchers is to estimate VARs over moving windows of data. As time progresses, earlier observations are discarded in favour of more recent data, and model coefficients are re-estimated.

too insensitive to recent inflation, while those from models that impose unit root restrictions on inflation tend to be excessively sensitive to recent inflation. Consequently, this paper follows Kozicki and Tinsley (2001a, b) by using a shifting-endpoint model to approximate the implicit forecasting model for inflation that underlies survey expectations.

An advantage of the shifting-endpoint AR specification is that it can capture the implications of structural change that lead to shifting long-horizon expectations. Moreover, since the model has relatively few parameters, it is less likely to overfit the data than more complicated time-series specifications. This section describes the basic linear shifting-endpoint AR model as it will be applied to inflation. However, because survey data correspond more closely to an average of forecast inflation over multiple months (and years), the implied relationship between historical data and survey data is non-linear in AR parameters of the shifting-endpoint model. This relationship is also derived. Finally, econometric approaches to deal with the unobserved endpoint and missing observations are reviewed.

2.1 A shifting-endpoint AR model

A standard autoregressive model describing the evolution of inflation (π_t) is:

$$\pi_{t+1} = \alpha(L)\pi_t + (1 - \alpha(1))\mu + \epsilon_{t+1}, \tag{1}$$

where $\alpha(L) \equiv \alpha_1 + \alpha_2 L + \cdots + \alpha_p L^{p-1}$ is a polynomial in the lag operater L, defined by $L\pi_t \equiv \pi_{t-1}$, and ϵ_t is an innovation, typically assumed to be independent Normal with mean zero. Standard models of inflation assume that inflation is either I(0), implying that all roots of $\alpha(L)$ lie outside the unit circle, or I(1), implying that one root lies on the unit circle and that remaining roots lie outside the unit circle. In the unit root case, $\alpha(1) = 1$ and the endpoint will be a moving average of order p of inflation. By contrast, if all roots of $\alpha(L)$ lie outside the unit circle, then π_t will revert to the endpoint, or mean (μ), in the long run; i.e., $\lim_{k-\infty} E_{t-1} y_{t+k} = \mu$. The forecasting model can be represented conveniently in companion form:

$$\pi_{t+1} = \iota_1' z_{t+1} = \iota_1' C z_t + \iota_1'(I - C)\iota\mu + \iota_1'\iota_1\epsilon_{t+1}, \tag{2}$$

where $z_t \equiv [\pi_t \quad \ldots \quad \pi_{t-p+1}]'$. I is a $p \times p$ identity matrix, ι is a $p \times 1$ vector of ones, ι_1 is a $p \times 1$ vector with a one in the first element and remaining elements zero, and

$$C \equiv \begin{bmatrix} \alpha_1 & \alpha_2 & \ldots & \alpha_{p-1} & \alpha_p \\ 1 & 0 & \ldots & 0 & 0 \\ 0 & 1 & \ldots & 0 & 0 \\ & & \ddots & & \vdots \\ 0 & 0 & \ldots & 1 & 0 \end{bmatrix}.$$

A simple approach to introduce limited forms of structural change into time-series systems is to admit shifts in the endpoints of the variables:

$$\lim_{k \to \infty} E_I \pi_{t+k} = \mu_\infty^{(I)}, \tag{3}$$

where I indexes the information set on which expectations are conditioned.[8] The general inflation-forecasting model with shifting endpoints can be represented as

$$\pi_{t+1} = \iota_1' z_{t+1} = \iota_1' C z_t + \iota_1'(I - C)\iota \mu_\infty^{(t)} + \iota_1' \iota_1 \epsilon_{t+1}. \tag{4}$$

As before, if all roots of $\alpha(L)$ lie outside the unit circle, then conditional long-horizon forecasts of π will revert to the endpoints $\mu_\infty^{(t)}$. Unlike the standard model, these endpoints may shift according to information and beliefs at the time the forecast is made. Intuitively, because the inflation endpoint is the conditional long-horizon forecast of inflation generated by the model, in a model of private sector expectations it can be thought of as the private sector perception of the inflation target.

The endpoint represents the level at which inflation forecasts are expected to eventually converge, conditional on a given information set. If survey participants could forecast future changes to their perceptions of the level at which inflation would stabilize, then such changes would be immediately incorporated. Consequently, changes in the endpoint should not be forecastable. This property is captured by assuming that the endpoint evolves according to a random walk:

$$\mu_\infty^{(t+1)} = \mu_\infty^{(t)} + v_{t+1}. \tag{5}$$

[8] Evidence of shifts in the mean of inflation are provided by Garcia and Perron (1996). They model inflation using a Markov switching specification with three states. As in their specification, parameters governing the speed of adjustment to long-run equilibrium (C) are assumed to be constant in the current implementation, even with shifts in the description of long-run equilibrium. The implications and relevance of other generalizations to the forecasting system are left for future research.

More details on the properties of ϵ_t and v_t will be provided with the state-space description of the model in section 2.3.

The shifting-endpoint specification described above is a generalization of the local-level model of Harvey (1989) and a version of the unobserved components model discussed by Watson (1986).[9] In particular, for $\theta^c(L) \equiv (1 - \alpha(L)L)^{-1}$ and $\tau_{t+1} \equiv \mu_\infty^{(t)}$, the shifting-endpoint specification can be rewritten as:

$$\pi_{t+1} = \tau_{t+1} + c_{t+1}, \tag{6}$$

$$\tau_{t+1} = \tau_t + v_{t+1}, \tag{7}$$

$$c_{t+1} = \theta^c(L)\epsilon_{t+1}. \tag{8}$$

In a recent study comparing several simple models, Stock and Watson (2005) find that a version of this specification with $\theta^c(L) = 1$ and time-varying estimates of the variances of ϵ_t and v_t performs remarkably well at forecasting inflation. The more general lag structure considered here is advantageous for capturing seasonality in data.

The shifting-endpoint specification shares features with other specifications proposed in the literature. For instance, the specification resembles the regressive-expectations model of Figlewski and Wachtel (1981). They express expected inflation as a weighted average of lagged inflation and long-run "normal" inflation, where the latter is defined as the rate towards which inflation is expected to regress. However, whereas Figlewski and Wachtel assume that the normal inflation rate is equal to a five-year moving average of inflation, here the shifting endpoint is treated as an unobserved component to be estimated. Caskey (1985) estimates a time-varying constant in a more general learning model of Livingston's 8-month inflation expectations. Caskey's learning model is a time-varying parameter model that includes a constant and several macroeconomic variables. He interprets a loose prior on the variance of the constant as evidence that the Livingston panel were willing to quickly revise their beliefs about the constant, and he concludes that Livingston inflation forecasts could be explained as the product of a learning process.

In other related work, Ang, Bekaert, and Wei (2005) find that a non-linear regime-switching model with two regimes (allowing both the mean and lag coefficients to switch) was a good forecasting specification for CPI inflation in the post-1995 period. They attribute this advantage to a reduction in the persistence of inflation at the end of the sample that can

[9]Unobserved components models are frequently used to model trend-cycle decompositions of real GDP (or GNP), as in Harvey (1985), Watson (1986), and Stock and Watson (1988).

be captured through a regime switch. By allowing the endpoint to follow a random walk, the shifting-endpoint model implicitly captures more than two regimes. Shifts of the endpoint capture structural change and absorb some of the persistence of inflation. Although AR parameters in C' are constant, lower persistence is captured with a decrease in the importance of endpoint movements relative to inflation deviations for explaining inflation dynamics at the end of the sample.

2.2 Approximating survey expectations with AR expectations

As outlined earlier, survey data provide timely information on perceived economic structural change. Because survey data on expectations include judgmental views as well as the output of econometric forecasting models, such data are likely to immediately reflect perceptions that there have been structural shifts in the economy. The consequences for inflation expectations of these perceptions of structural shifts can be extracted by linking the AR-based forecasting model to survey data on multiple-horizon expectations.

Survey forecasts report average inflation over multiple periods. Let $s_{t+k,t}$ denote the survey forecast of average expected inflation over the k periods ending in $t+k$, conditional on information available at t:

$$s_{t+k,t} = \frac{1}{k}\sum_{j=1}^{k} E_t^S \pi_{t+j}. \tag{9}$$

where E_t^S signifies that expectations are made by survey participants and conditional on information available at t.

Multi-step forecasts of inflation based on the shifting-endpoint AR model are:

$$\begin{aligned} E_t\pi_{t+j} &= \iota_1' E_t z_{t+j} \\ &= \iota_1' C^j z_t + \iota_1'(I - C^j)\iota\mu_\infty^{(t)}. \end{aligned} \tag{10}$$

and conditional forecasts of average inflation over the next k periods are:

$$(1/k)\sum_{j=1}^{k} E_t\pi_{t+j} = \iota_1'((1/k)\sum_{j=1}^{k} C^j)z_t + \iota_1'(I - ((1/k)\sum_{j=1}^{k} C^j))\iota\mu_\infty^{(t)}. \tag{11}$$

Assuming the average-inflation forecast from the shifting-endpoint AR model of inflation

provides an approximation of the survey forecast,

$$s_{t+k,t} = (1/k) \sum_{j=1}^{k} E_t \pi_{t+j} + \eta_{k,t}. \tag{12}$$

where $\eta_{k,t} = (1/k) \sum_{j=1}^{k} (E_t^S \pi_{t+k} - E_t \pi_{t+k})$ is an approximation error. The approximation error reflects differences between the implicit forecasting model of the survey participants and the shifting-endpoint AR model, and measurement error in the survey data, among other contributors. However, as both the survey data and the AR-based average-inflation forecast are conditioned on information in t, the approximation error does not reflect differences between actual inflation and predictions. For the same reason, there is no justification for expecting that approximation errors will be serially correlated. The latter point is in contrast to the difference between actual average inflation over k periods and k-period predictions, which will in general follow an MA(k-1).[10]

2.3 A state space model of the inflation endpoint

Estimates of parameters of the model and a time-series for the unobserved endpoint can be obtained by representing the model in state space format and using the Kalman filter to provide linear least squares predictions of the unobserved endpoint. State space representations, the Kalman filter, and approaches to estimating unobserved parameters are described in Harvey (1989) and Hamilton (1994).

In state space format, the endpoint is the unobserved state variable. As noted earlier, it is assumed to evolve according to a random walk:

$$\mu_\infty^{(t+1)} = \mu_\infty^{(t)} + v_{t+1}. \tag{13}$$

Innovations, v_t, are distributed Normal$(0, Q)$ with the mean square error matrix $Var_t(\mu_\infty^{(t+1)}) = P_{t+1|t}$.

Expressions for inflation and survey data constitute the measurement equations. Letting $k_1, k_2, \ldots k_n$ denote the various horizons for which the survey data are available, and defining $y_{t+1} = [\pi_{t+1} \;\; s_{t+k_1,t} \;\; s_{t+k_2,t} \;\; \ldots \;\; s_{t+k_n,t}]'$, the measurement equations are:

$$y_t = A' z_{t-1} + H' \mu_\infty^{(t)} + w_t, \tag{14}$$

[10]Hansen and Hodrick (1980) propose a methodology for examining restrictions on a k-step-ahead forecasting equation.

where

$$A' = \begin{bmatrix} \iota_1' C \\ \iota_1'((1/k_1)\sum_{j=1}^{k_1} C^j) \\ \iota_1'((1/k_2)\sum_{j=1}^{k_2} C^j) \\ \vdots \\ \iota_1'((1/k_n)\sum_{j=1}^{k_n} C^j) \end{bmatrix},$$

$$H' = \begin{bmatrix} \iota_1'(I - C)\iota \\ \iota_1'(I - (1/k_1)\sum_{j=1}^{k_1} C^j)\iota \\ \iota_1'(I - (1/k_2)\sum_{j=1}^{k_2} C^j)\iota \\ \vdots \\ \iota_1'(I - (1/k_n)\sum_{j=1}^{k_n} C^j)\iota \end{bmatrix}, \tag{15}$$

and $w_t = [\epsilon_{t+1} \;\; \eta_{k_1,t} \;\; \eta_{k_2,t} \;\; \dots \;\; \eta_{k_n,t}]'$ is distributed as Normal$(0, R)$, with v_t and w_t being independent of each other. The system described in (14) and (15) imposes the cross-equations restrictions necessary to ensure that the survey forecasts incorporate model-consistent expectations.

The structure of the covariance matrix, R, depends on the assumed relationships between inflation equation residuals (ϵ_{t+1}) and survey measurement errors ($\eta_{k,t}$), the assumed relationships between measurement errors of surveys of different horizons, and variances. Results are presented for the case of R diagonal with the variances of the measurement errors assumed to be the same for any choice of k_i, but with the variance of ϵ_{t+1} allowed to be different from the variance of measurement errors.[11]

Maximum likelihood estimation is described in Harvey (1989) and Hamilton (1994). Under normality of v_t and w_t, the log-likelihood function can be constructed using the Kalman filter. With starting values for the unobserved state and its mean square error, maximum likelihood techniques can be used to estimate parameters in A, H, Q, and R.

Basic intuition for the model follows from an examination of the data, shown in Figure 1. Notice that the data are generally ordered with inflation closest to the 8-month survey,

[11]Since the measurement equations have the state variable (i.e., the inflation endpoint or the perceived inflation target) as an explanatory variable, and since the forecasting model is expressed in deviation from the endpoint format, it is helpful to think of the measurement error as the measurement error in the deviation of the k-step-ahead forecast from the endpoint. Using this intuition, the measurement-error variance goes to zero as the horizon, k, increases; i.e., $lim_{k\to\infty} var(u_{kt}) = 0$. Although this is only an infinite horizon property, we tried various ways of approximating a tapering of the measurement-error variance over the available survey horizons but none were successful in obtaining stable or significant estimates.

followed by the 14-month survey and then the 10-year survey. This is exactly the ordering to expect if forecasts—and, hence, average forecasts, as in the case of survey expectations—are weighted averages of inflation and the inflation endpoint.[12] In particular, if all roots of $\alpha(L)$ lie outside the unit circle, then as j increases, the matrix C^j approaches a matrix of zeros. Thus, the larger the k_j, the smaller the weight (in A') on z_t and the higher the weight (in H') on $\mu_\infty^{(t)}$ in the corresponding measurement equation in (14). With respect to Figure 1, survey observations should be bounded by inflation on one side and the unobserved endpoint on the other, with shorter-horizon (i.e., smaller k_j) expectations closer to inflation and longer-horizon expectations closer to the endpoint.

This intuition reveals an important empirical advantage of using survey data to help estimate the endpoint. Deviations of predictions of survey expectations from actual survey expectations with large k_j will receive more weight when updating estimates of $\mu_\infty^{(t)}$ than those for small k_j. This is evident from the expression that describes Kalman updates of predictions of the state variable:

$$E_t\mu_\infty^{(t+1)} = E_{t-1}\mu_\infty^{(t)} + K_t(y_t - A'z_{t-1} - H'E_{t-1}\mu_\infty^{(t)}). \tag{16}$$

where K_t is the Kalman gain and is defined according to:

$$K_t = P_{t|t-1}H(H'P_{t|t-1}H + R)^{-1}. \tag{17}$$

All else equal, the matrix H embedded in K_t implies that deviations of long-horizon survey expectations from model predictions will obtain a weight close to one, while the weight on deviations of one-step-ahead inflation predictions from actual inflation will be much closer to zero. Since the model is expressed in a format where expectations converge to μ_∞ with horizon, this is exactly what one would want. Long-horizon expectations should provide more information about the limit of expectations (the endpoint) and, consequently, should receive more weight in estimating the endpoint.

2.4 Dealing with missing observations

One drawback of the Livingston Survey data is that they are available less frequently and for a shorter horizon than are inflation data.[13] One option would be to use observations

[12]The intuition may be clearer in the case of an AR(1) model of inflation where A and H are vectors with an ith entry of A equal to $w_i \equiv \alpha(1-\alpha^{k_i})/(k_i(1-\alpha))$. In this case, $s_{t+k_i,t} = w_i\pi_t + (1-w_i)\mu_\infty^{(t)}$, with $\lim_{k_i\to\infty} w_i = 0$.

[13]The Livingston Survey data are described in more detail in the next section.

for t only when data are available for every component of y_t. However, this would result in an extremely limited dataset, since long-horizon survey data are available only since the early 1990s. An alternative would be to drop observations for the long-horizon survey data, and include observations with shorter-horizon survey expectations and inflation. While this would expand the set of available observations considerably, analysis would still be limited to only two observations per year.

The approach taken in the next section is to use all available data starting in 1955. Using this approach, monthly observations are available for every year for the inflation measurement equation, semi-annual observations are available every year for the measurement equations of two relatively short-horizon survey expectation series, and semi-annual observations are available since 1991, with one observation from 1990.[14]

The methodology outlined in Harvey (1989, 144) is used to deal with missing observations. In particular, the model just described is transformed into a system with measurement equations for $y_t^* = W_t y_t$, where W_t is a matrix that selects those elements of y_t for which observations are available. In the description of the measurement equations, $A_t^{*\prime} = W_t A'$, $H_t^{*\prime} = W_t H'$, and $R_t = W_t R W_t'$, respectively, replace A', H', and R.

3. Empirical Results

3.1 Data

Survey data on short-horizon CPI expectations are taken from the Livingston Survey. This survey is conducted twice per year, in June and December.[15] Participants are asked to give 6-month and 12-month forecasts of the CPI level. However, because CPI data are released with a lag, the recommendation of Carlson (1977) is followed and it is assumed that, when making their forecasts, economists had access to CPI data through April and October, respectively. Thus, the survey data are treated as 8-month and 14-month forecasts of the CPI level. While informational assumptions may differ across survey participants, Carlson (1977) reports that this assumption is likely consistent with the practice of the majority of those surveyed.

[14]Results from estimations that exclude long-horizon survey expectations entirely, or that use only semi-annual observations of inflation and shorter-horizon survey expectations, were used to check the robustness of the results.

[15]Documentation describing the Livingston Survey data is available on the Federal Reserve Bank of Philadelphia website at http://www.phil.frb.org. Croushore (1997) provides a description of the survey and its history.

A complication that arises when trying to use the survey data is that in a few instances since the start of the survey, the CPI has been rebased to 100 and rounded, but the survey levels have not been rebased. To minimize distortions that rounding and rebasing introduce, the alternative base-year CPI published by the Bureau of Labor Statistics (rebased with 1967=100) is used for the empirical analysis and both survey data and price-level data are converted to inflation rates. As reported by Kozicki and Hoffman (2004), distortions associated with rounding are considerably smaller in the alternative base-year CPI, and inflation rates will be comparable even if the index levels of the actual and survey series are not scaled to the same base year.[16]

Another feature of the Livingston Survey data is that the CPI being forecast is not a seasonally adjusted series. For this reason, an AR(13) specification is used. Specifications with fewer lags were also considered, but tended to generate excessively volatile near-term forecasts.[17]

3.2 Results

Motivation for the choice of the shifting-endpoint specification is based on the failure of constant-endpoint and unit root models of inflation to match survey data in a different set-up (Kozicki and Tinsley 1998, 2001a, b). Since those studies do not use survey data during estimation, and their conclusions are based on a different survey, the performance of these alternatives might be better in the current application. For this reason, results from these specifications are included for comparison.

The constant-endpoint AR specification for inflation is given in (1). The unit root specification is a restricted version of (1) where $\alpha(1) = 1$ has been imposed.[18] A transition equation describing the evolution of the endpoint is not required for either of these variants. Thus,

[16]CPI data are generally not revised, so the only differences between inflation calculated using the alternative base-year CPI and real-time data are due to rounding that may occur during rebasing. In preliminary work on semi-annual data, real-time CPI data were used and results similar to those reported in the paper were obtained.

[17]In preliminary work, autoregressive specifications with seasonal dummies were less successful at capturing the seasonality. Moreover, coefficients on seasonal dummies tended to be insignificantly different from zero.

[18]The shifting-endpoint specification as implemented in this paper admits a unit root in inflation. Consequently, it can be seen as a restricted version of a unit root specification. The advantage of the shifting-endpoint specification is that it provides a parsimonious alternative to an unrestricted unit root specification with long lags. For instance, Jorda (2005) argues that very long lags are needed to match impulse responses in inflation. That said, the unit root in the endpoint is an ex post interpretation of a parsimonious local approximation. From the real-time perspective of agents, the series is conditionally stable about the current endpoint. As seen by the sum of AR coefficients in Table 1, the largest root about the shifting endpoint is about half the size of the root about a constant endpoint.

parameters in A, H, and R (and μ in the constant-endpoint case) are estimated by applying maximum likelihood to the measurement equations summarized in (14).

To proceed with maximum likelihood estimation of the shifting-endpoint specification, starting values for the endpoint and its mean square error are required. Given the random-walk transition equation for the shifting endpoint, a diffuse prior is assumed. In particular, the mean square error is set to 1000 and the mean is set to 2.5 per cent (the value of μ estimated in the constant-endpoint variant).

Results using data from 1955 through April 2005 are summarized in Table 1.[19] In many respects, the models are similar. Point estimates of individual autoregressive parameters (α_i) are similar: estimated coefficients on the first lag are all slightly larger than 0.3, and all models capture seasonality in the data with statistically significant estimates of the coefficient on the twelfth lag close to 0.2. In addition, standard errors of the measurement equation for inflation differ by less than 0.01 percentage point, suggesting that the three specifications explain the behaviour of inflation equally well at one-month horizons.

The key difference between the three specifications is that persistence as measured by the sum of autoregressive coefficients is lower in the shifting-endpoint specification than in the constant- or moving-average- (MA-) endpoint specifications. This result is consistent with Kozicki and Tinsley (2002), who report a notable decline in the sum of AR coefficients after allowing for a shifting endpoint, and with Kozicki and Tinsley (2001b), who find that unit root tests on the deviation of inflation from an estimated inflation endpoint are rejected and that those on inflation are not. In an extension to multiple countries, Levin and Piger (2004) confirm that inflation persistence decreases after accounting for mean shifts. The intuition behind these results is that some of the persistence in inflation is absorbed into low-frequency movements of the shifting endpoint that anchor long-horizon inflation expectations. In Figure 2, movements of the smoothed estimate of the shifting endpoint lag low-frequency movements in CPI inflation (expressed in the figure as inflation over the prior 12 months).

[19] While use of the longest possible sample (1946 is the first year for which 8- and 14-month surveys are available) is desired, three factors motivate consideration of a somewhat shorter sample. First, as noted by Carlson (1977), Livingston tends to adjust survey data with the release of inflation data for months prior to the survey date. Such adjustments in the first part of the survey history may distort the data relative to more recent observations. Second, distortions owing to rounding and rebasing of CPI data are larger for earlier observations. Finally, inflation itself appears to be generated by a different process in the years following WWII—inflation is more variable and the duration of lower-frequency fluctuations is shorter. The choice of 1955 as a starting observation reflects a compromise, and a robustness check suggests that similar results are obtained for shorter samples.

The estimated specifications can be used to construct term structures of expected inflation: i.e., profiles of expected inflation over different forecast horizons. Model estimation provides monthly observations of the shifting endpoint. This series, combined with the estimated model parameters and monthly inflation data, can be used to construct monthly forecasts of inflation at any horizon using expression (10), and predictions of average inflation over any horizon using expression (11). Thus, although available survey data are limited to *semi-annual* observations on only *three* horizons, the model can be used to construct *monthly* predictions at *any* horizon.

The results of such an exercise are presented in Figure 3a for the shifting-endpoint specification, and in Figures 3b and c, respectively, for the constant-endpoint and unit root specifications.[20] The profiles for the shifting-endpoint specification in Figure 3a show relatively fast reversion of inflation expectations to the endpoint as the forecast horizon increases. But, owing to time variation in the endpoint, predictions of long-horizon inflation expectations incorporate considerable variation over history. By contrast, higher estimated persistence implies more gradual mean reversion and sluggish adjustments of near-term inflation expectations in the constant-endpoint specification of Figure 3b. However, as forecasts revert to a constant in this specification, long-horizon inflation expectations exhibit relatively little variation. Finally, the unit root restriction in the third specification implies that forecasts at all horizons remain close to recent inflation, as shown in Figure 3c.

A second important difference between the specifications is in their ability to match survey expectations. The standard error of the measurement equations for the survey data is considerably smaller for the shifting-endpoint specification than for the other two specifications. This result provides an early indication that the shifting-endpoint specification comes closer to fitting survey data than the others.

Survey expectations and predictions based on the three specifications are shown in Figure 4 for the 8-month forecast horizon, and in Figure 5 for the 10-year horizon.[21] In both cases, the shifting-endpoint prediction tracks the survey data quite closely. However, in Figure 4, both the constant-endpoint and unit root specifications generate predictions of 8-month inflation expectations that are more volatile than actual inflation.

[20]Considerable month-to-month volatility in inflation implies a fair degree of month-to-month volatility in near-term inflation predictions; consequently, to make the figures easier to examine, only two profiles are shown per year. Likewise, predictions are shown only for every third horizon: i.e., for horizons 1, 4, 7, 10, 13, ... months.

[21]A figure showing results for the 14-month horizon has been excluded because the results are visually similar to those for the 8-month horizon.

In Figure 5, the shifting-endpoint specification generates 10-year inflation predictions that appear to provide a compromise between predictions based on the other two specifications. In particular, the prediction from the constant-endpoint specification exhibits relatively little variation and appears strongly anchored to 2.5 per cent over most of the sample. At the other extreme, the unit root specification predicts considerable volatility and, owing to the unit root restriction, follows actual inflation closely.

Table 2 provides formal evidence on the superior ability of the shifting-endpoint specification to match available survey data. Entries are root mean squared deviations (RMSD) between survey data and model-based predictions of multi-period inflation forecasts. What is interesting about this comparison is that the shifting-endpoint specification clearly dominates the other specifications, even though all three specifications are fit to survey data and inflation, and the ability of each to fit inflation is similar. Thus, using survey data during estimation and having a good model of inflation are jointly *not* sufficient to generate good proxies for expected inflation.

Evidence on the ability of the models to fit 10-year survey data since 1990 might not be seen as very strong. After all, 8- and 14-month survey expectations exhibit considerably more volatility prior to 1990 than afterwards, and 10-year Livingston expectations are not available in the earlier part of the sample. Consequently, the limited history of 10-year Livingston data might be seen as an impediment to the evaluation of the model. Moreover, in order for model predictions of long-horizon inflation expectations to be taken as reasonable proxies for survey data, additional evidence on fit prior to 1991 would be valuable. Such evidence is provided in Figure 6 and in the final row of Table 2.

Figure 6 compares the three predictions of 10-year inflation expectations from the shifting-endpoint specification with the limited Livingston Survey data and with spliced survey data on long-horizon inflation expectations. The spliced survey data are taken from the *Blue Chip Economic Indicators* (available twice per year) through March 1991, and from the *Survey of Professional Forecasters* from November 1991 through to the end of the sample (available quarterly). Although the Livingston Survey data are used during estimation, the spliced survey data are not. Thus, the spliced survey data provide an external check on the validity of the predictions from the shifting-endpoint specification. The surveys track each other closely when both are available, suggesting that the views summarized by the two surveys are well aligned. In fact, the shifting-endpoint predictions track the path of the spliced survey observations quite closely and fluctuations in the two are synchronized. That said, there is weak evidence that long-horizon predictions are a little too sensitive to recent movements

in inflation. Relative to the spliced survey data, predictions are somewhat high prior to the Volcker disinflation, and somewhat low afterwards.[22]

Nevertheless, 10-year inflation predictions of the shifting-endpoint specification clearly dominate the predictions of the constant-endpoint and unit root specifications in their ability to match survey expectations (Table 2). RMSDs between available spliced survey expectations and the predictions are 75 per cent larger for the constant-endpoint specification, and over twice as large for the unit root specification. Thus, this comparison provides additional evidence that the shifting-endpoint predictions are reasonable, including in the period prior to 1990.

While the analysis discussed so far is conditioned on choices regarding sample period, autoregressive lag length, and inclusion of very limited 10-year survey data, further investigation provides evidence that the results are remarkably robust. Table 3 shows that estimation results are similar for three different sample periods. Estimates of persistence are in the range of 0.45, with the largest AR coefficient applying to the first lag on inflation, and standard errors on the innovation to the state variable are close to 0.23. Although the estimated first autoregressive coefficient is somewhat larger for the shortest sample than for the other two, the implications are largely unwound by more negative second and third autoregressive coefficients.

Table 4 compares results from the baseline shifting-endpoint specification already discussed to a variant that excludes the survey data on 10-year inflation expectations. Parameter estimates, including the sum of AR coefficients, are very close. In addition to establishing robustness, these results suggest that even relatively short-horizon expectations provide considerable information on long-horizon perceptions.

A final check on the robustness of the results is provided by comparing the shifting endpoint that is estimated to anchor Livingston Survey expectations with comparable constructions from other studies. Figure 7 contains an estimate of the "normal" inflation rate (Figlewski and Wachtel 1981), an inflation endpoint based on an adaptive-learning model (Kozicki and Tinsley 2001b), an inflation endpoint based on a changepoint-learning model

[22]This might be due to distortions in the 8-month and 14-month survey data that resulted from adjustments to the raw survey data made by Livingston. As noted by Carlson (1977), when new data were released between the time of the survey being conducted and the time of its results being published, Livingston sometimes adjusted raw survey data in the direction of surprises in the data. Alternatively, the assumption in the model that the AR parameters were constant over the entire sample may be overly restrictive. Cogley and Sargent (2005) find evidence of time variation in the persistence of inflation even when allowing for a shifting mean.

(Kozicki and Tinsley 2001b), and a VAR-based perceived inflation target (Kozicki and Tinsley 2005a). The similarities of these five series is striking, particularly given the differences in the underlying data and methodologies. All of the series shown in Figure 7 move gradually, with general increases in the 1960s and 1970s, and decreases in the 1980s and (to a lesser extent) 1990s. However, of the estimates shown, only the shifting endpoint constructed in this paper makes use of survey data during estimation. The normal inflation rate and adaptive-learning-model estimate of the inflation rate are both moving averages of past inflation—the former equally weights inflation (over the prior five years), whereas the latter uses weights that decline geometrically. The changepoint-learning model approximates real-time learning using breakpoint tests with expanding samples to detect mean shifts in an AR model of inflation. The VAR-based perceived inflation target is an unobserved component that enters into the central tendencies of inflation and nominal interest rates, and is assumed to shift inversely with unanticipated policy shocks.

4. Heterogeneous Perceptions of Inflation Targets

Although monetary policy in the United States is conducted without announced numerical targets for inflation, policy decisions are designed with inflation objectives in mind. Likewise, nominal-debt contracts, wage- and price-setting behaviour, and other economic decisions by households and firms are influenced by inflation expectations, which are anchored by private perceptions of the central bank's inflation target. In the absence of an announced numerical inflation goal and full information, private and central bank perceptions of the effective inflation target may diverge.

The shifting endpoint estimated in the previous section provides a measure of private sector perceptions of the implicit inflation goal of monetary policy. These private sector perceptions can be compared with estimates of central bank perceptions to assess policy credibility. Kozicki and Tinsley (2005b) estimate the effective target of monetary policy using real-time Federal Reserve Board staff forecast data. Alternative estimates of the effective inflation target from an unobserved components model used with retrospective data are provided by Kozicki and Tinsley (2005a). To the extent that low-frequency movements in actual inflation may predominantly reflect the effective goal of policy, the estimate of "core inflation" in Cogley and Sargent (2005) may also proxy for the effective target of policy.

Figure 8 illustrates divergences between the effective inflation target and private sector perceptions. By all three measures of the effective target, policy actions through the 1970s

were *as if* the central bank was willing to achieve inflation of roughly 6 to 7 per cent. By contrast, the private sector was slow to adjust their views, and their perceptions of the inflation goal increased only gradually, from about 3 per cent in 1970 to about 7 per cent by the end of the decade.

The opposite outcome was observed in the 1980s. All three measures of the effective target exhibited a rapid decline near the end of 1979. However, private sector perceptions adjusted much more slowly. Gaps between private sector perceptions and the central bank's effective target provide evidence that the Volcker disinflation was not initially viewed as being fully credible.

These results on policy credibility are reinforced by re-examining the preferred estimates of the term structures of expected inflation (Figure 3a). In particular, the term structures facilitate a comparison of realized inflation with constructed forecasts of long-horizon inflation. The analysis shows that, following the Volcker disinflation, the term structure of expected inflation remained upward sloping, suggesting that market participants did not believe that policy would keep inflation rates at moderate levels. By contrast, throughout the 1990s, the term structure of expected inflation gradually flattened, suggesting that, under Greenspan, a monetary policy goal of low and stable inflation gained credibility.

An important feature evident in Figure 8 is the lag in low-frequency movements of private sector perceptions compared with the effective inflation-target series. A similar lag is evident between actual inflation and private sector perceptions (Figure 2). This phase shift is essential for explaining the behaviour of surveys and also the behaviour of long-term bond rates.[23] Time variation of coefficients, by itself, is not enough to capture the lags involved in real-time learning. For instance, the Cogley and Sargent (2005) VAR admits random walks in intercepts and slopes (although the latter are constrained to yield a stable VAR), yet their core inflation measure does not capture the phase shift in endpoints displayed in surveys and financial forecasts. Using data from surveys or financial forecasts implicit in bond yields during estimation enables the shifting-endpoint specification to capture the phase shift.

5. Concluding Comments

The paper has described and implemented a methodology for constructing a 50-year monthly term structure of expected inflation that is consistent with infrequent observations of survey

[23] See the discussion in Kozicki and Tinsley (1998, 2001a, b).

data. A shifting-endpoint AR model of inflation fits inflation to an extent that is comparable with more commonly implemented AR models with constant endpoints or unit root constraints imposed. However, even when survey data are used during estimation, the latter two models are incapable of matching the profiles of survey data on expected inflation. Forecasts from constant-endpoint models are too volatile at short forecast horizons and too flat at long horizons. Forecasts from unit root specifications are excessively volatile at all horizons. An important lesson from this analysis is that models that fit inflation well may not provide good proxies for expected inflation, even if survey data are used during estimation.

The analysis also suggests that survey data on near-term inflation expectations contain considerable information about long-horizon views. Model estimates are similar regardless of whether 10-year inflation expectations are included during estimation.

The model provides an estimate of private sector perceptions of the effective inflation goal of monetary policy. Divergences between private sector perceptions and estimates of the effective inflation target from other studies provide evidence on historical levels of monetary policy credibility. Indeed, the paper finds strong evidence of heterogeneous perceptions of inflation targets for U.S. monetary policy.

References

Ang, A., G. Bekaert, and M. Wei. 2005. "Do Macro Variables, Asset Markets or Surveys Forecast Inflation Better?" NBER Working Paper No. 11538.

Breedon, F.J. and J.S. Chadha. 1997. "The Information Content of the Inflation Term Structure." Bank of England Working Paper No. 75.

Carlson, J. 1977. "A Study of Price Forecasts." *Annals of Economic and Social Measurement* 6: 27–56.

Caskey, J. 1985. "Modeling the Formation of Price Expectations: A Bayesian Approach." *American Economic Review* 75: 768–76.

Christensen, I., F. Dion, and C. Reid. 2004. "Real Return Bonds, Inflation Expectations, and the Break-Even Inflation Rate." Bank of Canada Working Paper No. 2004-43.

Cogley, T. and T.J. Sargent. 2005. "Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII US." *Review of Economic Dynamics* 8(2): 262–302.

Côté, A., J. Jacob, J. Nelmes, and M. Whittingham. 1996. "Inflation Expectations and Real Return Bonds." *Bank of Canada Review* (Summer): 41–53.

Croushore, D. 1997. "The Livingston Survey: Still Useful After All These Years." Federal Reserve Bank of Philadelphia *Business Review* (March/April): 1–12.

Figlewski, S. and P. Wachtel. 1981. "The Formation of Inflationary Expectations." *Review of Economics and Statistics* 63: 1–10.

Garcia, R. and P. Perron. 1996. "An Analysis of the Real Interest Rate Under Regime Shifts." *Review of Economics and Statistics* 78: 111–25.

Hamilton, J.D. 1994. *Time Series Analysis*. Princeton, NJ: Princeton University Press.

Hansen, L.P. and R.J. Hodrick. 1980. "Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis." *Journal of Political Economy* 88: 829–53.

Harvey, A.C. 1985. "Trends and Cycles in Macroeconomic Time Series." *Journal of Business and Economic Statistics* 3: 216–27.

———. 1989. *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge: Cambridge University Press.

Harvey, C.R. 1988. "The Real Term Structure and Consumption Growth." *Journal of Financial Economics* 22: 305–33.

Jordà, O. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review* 95(1): 161–82.

Kozicki, S. and B. Hoffman. 2004. "Rounding Error: A Distorting Influence on Index Data." *Journal of Money, Credit, and Banking* 36(3): Part 1, 319–38.

Kozicki, S. and P. Tinsley. 1998. "Moving Endpoints and the Internal Consistency of Agents' Ex Ante Forecasts." *Computational Economics* 11: 21–40.

———. 2001a. "Term Structure Views of Monetary Policy Under Alternative Models of Agent Expectations." *Journal of Economic Dynamics and Control* 25: 149–84.

———. 2001b. "Shifting Endpoints in the Term Structure of Interest Rates." *Journal of Monetary Economics* 47: 613–52.

———. 2002. "Alternative Sources of the Lag Dynamics of Inflation." In *Price Adjustment and Monetary Policy*, 3-47. Proceedings of a conference held by the Bank of Canada, November 2002. Ottawa: Bank of Canada.

———. 2005a. "Permanent and Transitory Policy Shocks in an Empirical Macro Model with Asymmetric Information." *Journal of Economic Dynamics and Control* 29: 1985–2015.

———. 2005b. "Perhaps the FOMC Did What It Said It Did: An Alternative Interpretation of the Great Inflation." Federal Reserve Bank of Kansas City Economic Research Department Working Paper No. RWP 05–04.

Laubach, T. and J.C. Williams. 2003. "Measuring the Natural Rate of Interest." *Review of Economics and Statistics* 84(4): 1063–70.

Levin, A.T. and J.M. Piger. 2004. "Is Inflation Persistence Intrinsic in Industrial Economies?" European Central Bank Working Paper No. 334.

McNees, S.K. 1986. "Forecasting Accuracy of Alternative Techniques: A Comparison of U.S. Macroeconomic Forecasts." *Journal of Business and Economic Statistics* 4: 5–15.

Reifschneider, D.L., D.J. Stockton, and D.W. Wilcox. 1997. "Econometric Models and the Monetary Policy Process." *Carnegie-Rochester Conference Series on Public Policy* 47: 1–37.

Roberts, J.M. 1995. "New Keynesian Economics and the Phillips Curve." *Journal of Money, Credit, and Banking* 27: 975–84.

———. 1997. "Is Inflation Sticky?" *Journal of Monetary Economics* 39: 173–96.

Shen, P. and J. Corning. 2001. "Can TIPS Help Identify Long-Term Inflation Expectations?" Federal Reserve Bank of Kansas City *Economic Review* (Fourth Quarter): 67–81.

Sims, C.A. 2002. "The Role of Models and Probabilities in the Monetary Policy Process." *Brookings Papers on Economic Activity* 2: 1–62.

Söderlind, P. and L.E. Svensson. 1997. "New Techniques to Extract Market Expectations from Financial Instruments." *Jounal of Monetary Economics* 40: 383–429.

Stock, J.H. and M.W. Watson. 1988. "Variable Trends in Economic Time Series." *Journal of Economic Perspectives* 2: 147–74.

———. 2005. "Has Inflation Become Harder to Forecast?" Unpublished manuscript.

Wallis, K.F. 1989. "Macroeconomic Forecasting: A Survey." *Economic Journal* 99(4): 28–61.

Wallis, K.F., P. Fisher, J. Longbottom, and J. Whitley. 1986. *Models of the UK Economy: A Third Review by the ESRC Macroeconomic Modelling Bureau.* Oxford: Oxford University Press.

———. 1987. *Models of the UK Economy: A Fourth Review by the ESRC Macroeconomic Modelling Bureau.* Oxford: Oxford University Press.

Watson, M. 1986. "Univariate Detrending Methods with Stochastic Trends." *Jounal of Monetary Economics* 18: 49–75.

Table 1: Estimation Results

	Shifting endpoint		Constant endpoint		MA endpoint	
Parameter	Estimate	SE*	Estimate	SE	Estimate	SE
α_1	.325	.036	.344	.038	.351	.039
α_2	.020	.040	.040	.042	.078	.042
α_3	-.073	.039	-.025	.040	-.026	.041
α_4	.001	.037	.050	.040	.082	.040
α_5	.049	.041	.046	.042	.052	.042
α_6	.012	.041	.003	.041	.012	.041
α_7	.056	.038	.078	.041	.087	.041
α_8	-.041	.041	-.003	.042	.026	.042
α_9	.013	.038	.029	.040	.012	.040
α_{10}	-.055	.038	.019	.041	.047	.040
α_{11}	.049	.039	.067	.042	.080	.042
α_{12}	.172	.038	.198	.041	.204	.041
α_{13}	-.085	.031	-.025	.035	-.008	
$\sum_i \alpha_i$	.445		.819		1.000	
μ			2.575	.153	-2.615	.736
$R_\pi^{1/2}$	2.732	.112	2.726	.079	2.716	.079
$R_{ss}^{1/2}$	.243	.016	.909	.044	1.047	.051
$Q^{1/2}$	.232	.021				

*SE = standard error.

All models are estimated using maximum likelihood with data starting in 1955. The shifting-endpoint specification employs Kalman filtering techniques to estimate the unobserved state variables (the perceived inflation target). The variance covariance matrix of the measurement equations is restricted to be diagonal during estimation, and variances of measurement equations for survey data are assumed to be the same. Results are presented for three AR (13) model specifications. The *shifting-endpoint* model has a shifting mean, estimated using a Kalman filter procedure; the *constant-endpoint* model is a standard unrestricted AR(13) process with a constant mean; and the *MA-endpoint* model is an AR(13) with a unit root restriction imposed (i.e., the sum of AR coefficients is constrained to equal one).

Table 2: Comparison of Fits to Survey Data

Forecast horizon (Survey)	Shifting endpoint	Constant endpoint	MA endpoint
8 month (Livingston)	0.22	0.94	1.35
14 month (Livingston)	0.14	0.93	1.39
10 year (Livingston)	0.25	0.65	0.68
10 year (Blue Chip)	0.40	1.29	1.39

This table contains root mean squared errors (RMSEs) constructed as the square root of the average squared deviation of inflation predictions from survey data over those observations for which survey data are available. The row labelled *10 year (Livingston)* uses 10-year inflation-expectations data from the Livingston Survey. These are the data that are used during estimation. The row labelled *10 year (Blue Chip)* uses 10-year inflation-expectation data from the *Blue Chip Economic Indicators* through March 1991, and from the *Survey of Professional Forecasters* from November 1991 through to the end of the sample (available quarterly). These data are not used during estimation. Inflation predictions are constructed over the reported horizon for three different time-series models of inflation. All three models are AR(13) specifications. The *shifting-endpoint* model has a shifting mean, estimated using a Kalman filter procedure; the *constant-endpoint* model is a standard unrestricted AR(13) process with a constant mean; and the *MA-endpoint* model is an AR(13) with a unit root restriction imposed (i.e., the sum of AR coefficients is constrained to equal one). Estimates of model parameters are provided in Table 1.

Table 3: Robustness of Results to Sample

	1955Q1 - 2005Q4		1965Q1 - 2005Q4		1975Q1 - 2005Q4	
Parameter	Estimate	SE*	Estimate	SE	Estimate	SE
α_1	.325	.036	.367	.040	.461	.046
α_2	.020	.040	-.029	.047	-.160	.057
α_3	-.073	.039	-.043	.042	-.011	.057
α_4	.001	.037	-.026	.042	-.047	.053
α_5	.049	.041	.087	.047	.127	.056
α_6	.012	.041	-.011	.047	-.118	.055
α_7	.056	.038	.101	.044	.184	.053
α_8	-.041	.041	-.068	.046	-.075	.054
α_9	.013	.038	.003	.042	.001	.051
α_{10}	-.055	.038	-.063	.043	-.075	.051
α_{11}	.049	.039	.080	.045	.117	.053
α_{12}	.172	.038	.156	.044	.123	.052
α_{13}	-.085	.031	-.071	.035	-.073	.038
$\sum_i \alpha_i$	.445		.482		.455	
$R_\pi^{1/2}$	2.732	.112	2.767	.090	2.587	.098
$R_{ss}^{1/2}$	.243	.016	.237	.017	.239	.019
$Q^{1/2}$	.233	.021	.257	.021	.235	.026

*SE = standard error.

Results are obtained using maximum likelihood estimation with Kalman filtering techniques to estimate the unobserved-state variables (the perceived inflation target). The variance covariance matrix of the measurement equations is restricted to be diagonal during estimation. and variances of measurement equations for survey data are assumed to be the same.

Table 4: Robustness of Results to Use of 10-Year Survey Data

	10-year survey used		10-year survey not used	
Parameter	Estimate	Standard error	Estimate	Standard error
α_1	.325	.036	.320	.037
α_2	.020	.040	.043	.040
α_3	-.073	.039	-.053	.039
α_4	.001	.037	.024	.038
α_5	.049	.041	.032	.041
α_6	.012	.041	-.017	.041
α_7	.056	.038	.027	.039
α_8	-.041	.041	-.028	.041
α_9	.013	.038	.038	.037
α_{10}	-.055	.038	-.022	.038
α_{11}	.049	.039	.037	.039
α_{12}	.172	.038	.139	.038
α_{13}	-.085	.031	-.088	.031
$\sum_i \alpha_i$	.445		.452	
$R_\pi^{1/2}$	2.732	.112	2.727	.079
$R_{ss}^{1/2}$	.243	.016	.203	.015
$Q^{1/2}$	.232	.021	.240	.021

All models are estimated with data starting in 1955. Results are obtained using maximum likelihood estimation with Kalman filtering techniques to estimate the unobserved-state variables (the perceived inflation target). The variance covariance matrix of the measurement equations is restricted to be diagonal during estimation, and variances of measurement equations for survey data are assumed to be the same.

Figure 1: Inflation and Survey Expectations

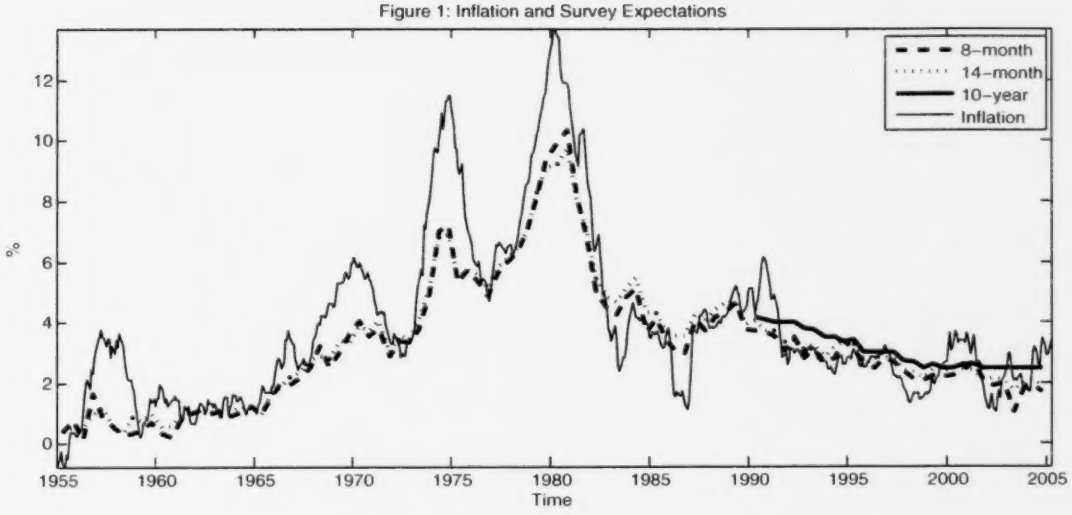

Figure 2: The shifting endpoint absorbs some inflation persistence

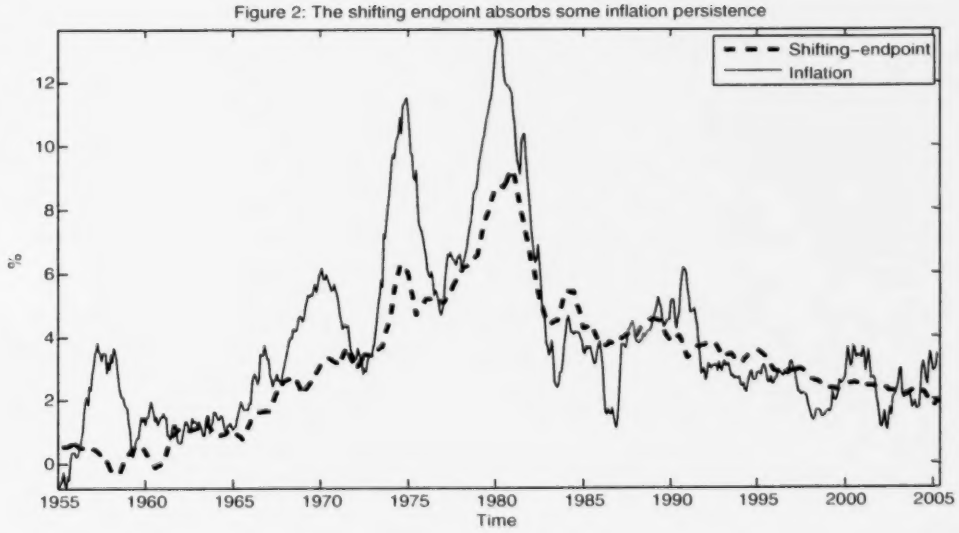

Figure 3a: The Term Structure of Expected Inflation
Shifting-endpoint specification

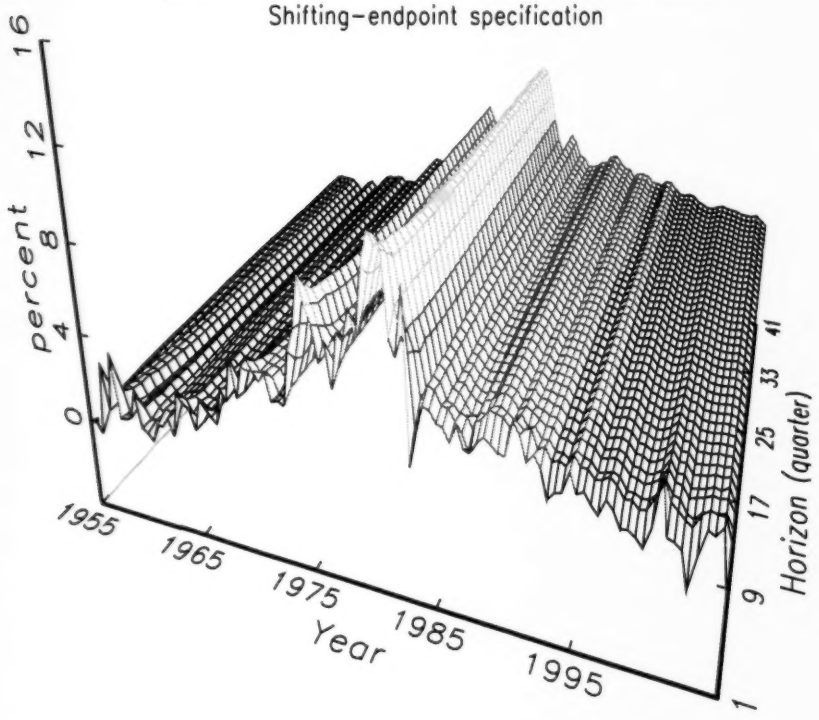

Figure 3b: The Term Structure of Expected Inflation
Constant endpoint specification

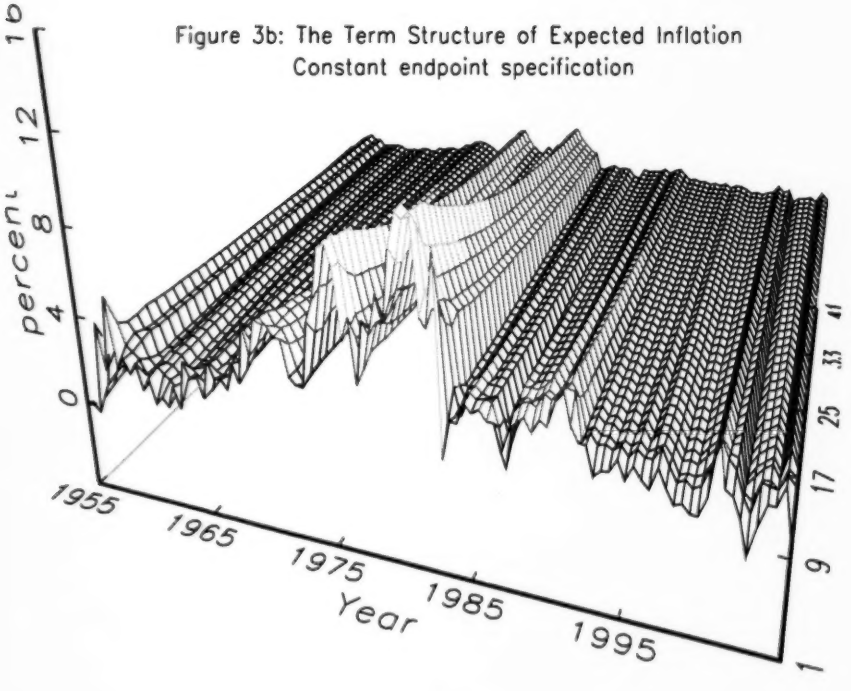

Figure 3c: The Term Structure of Expected Inflation
Unit root specification

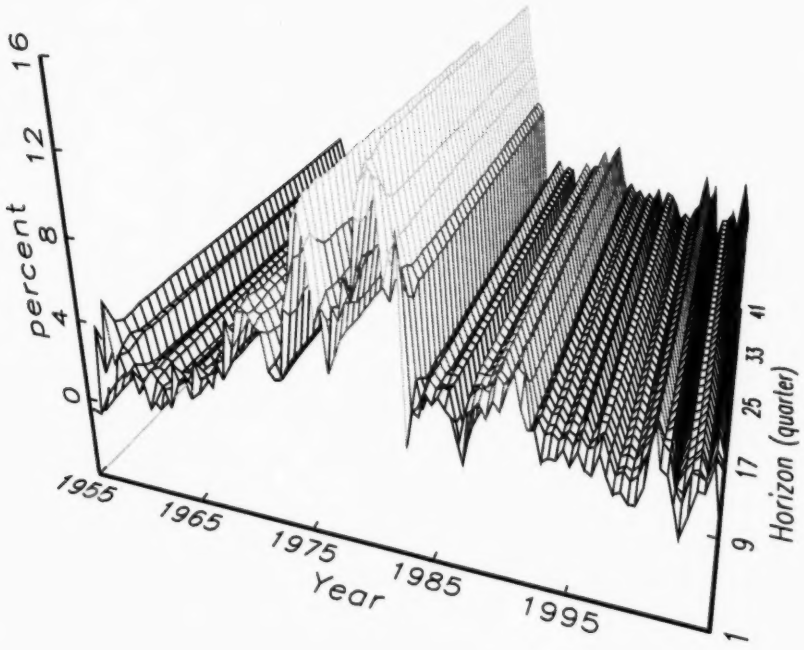

Figure 4: 8–month survey expectations and predictions

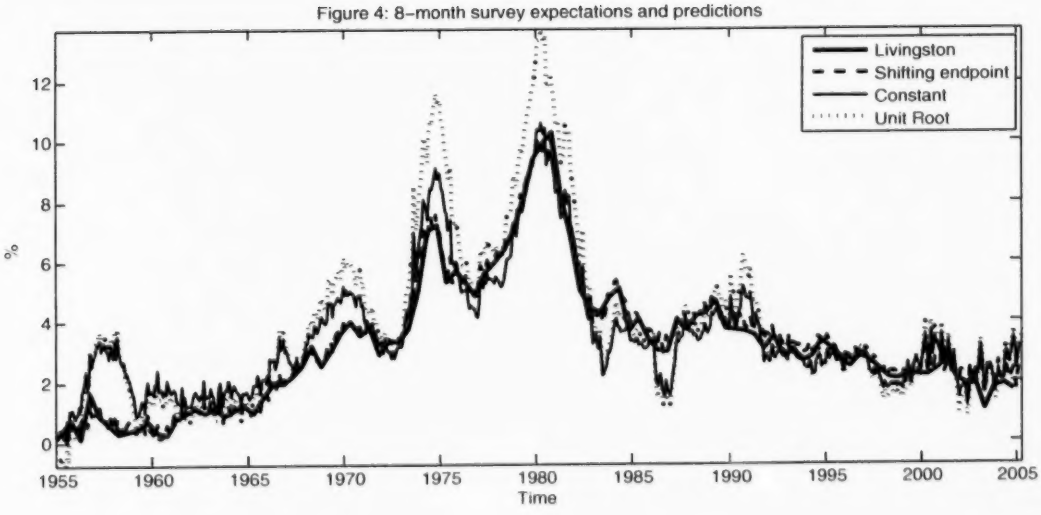

Figure 5: 10–year survey expectations and predictions

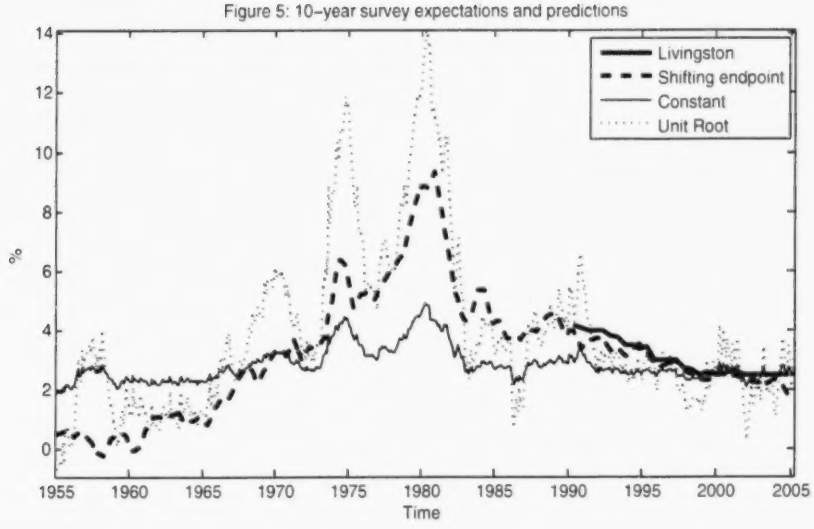

Figure 6: A Robustness check on predictions of 10–year expected inflation

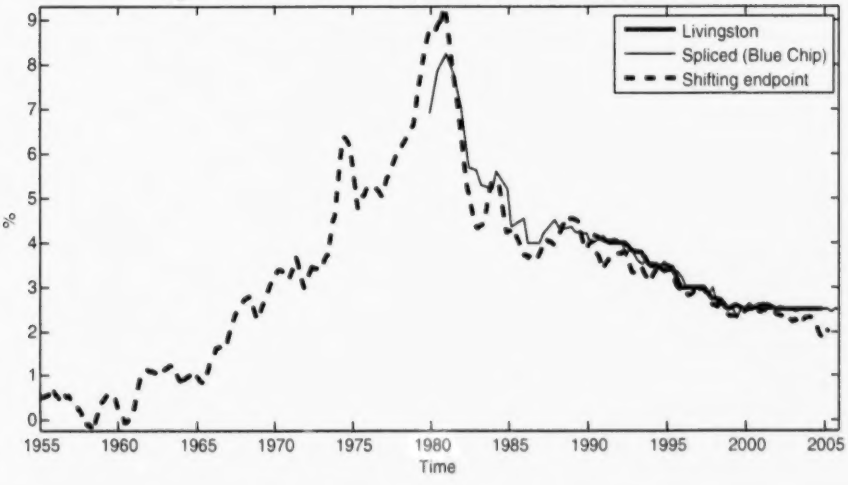

Figure 7: A comparison of alternative endpoint estimates

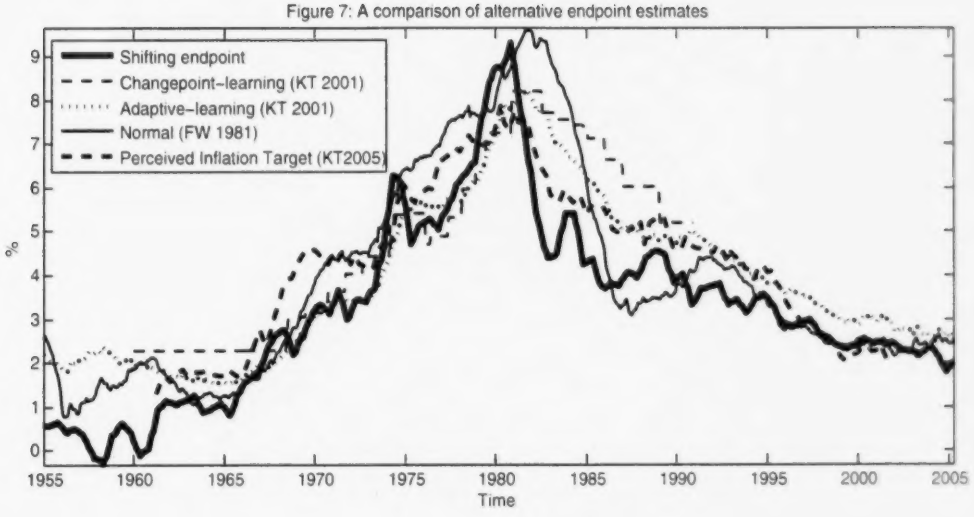

Figure 8: Heterogeneous perceptions

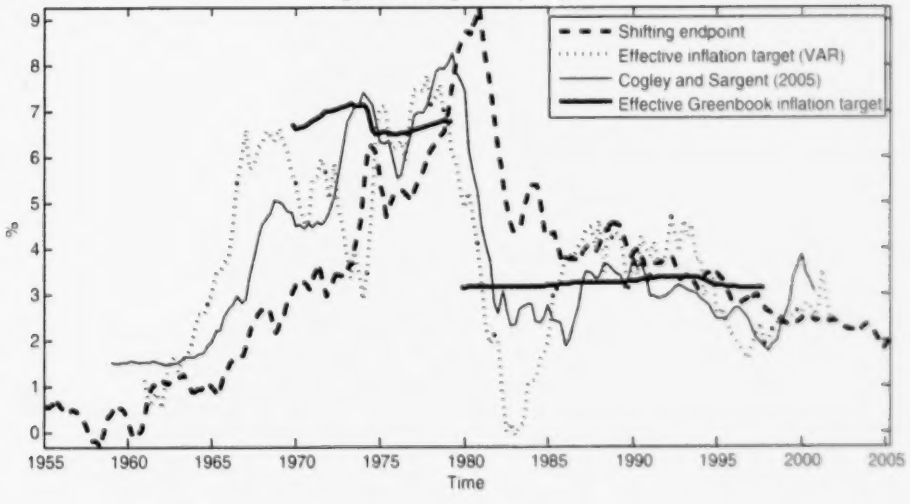

Bank of Canada Working Papers
Documents de travail de la Banque du Canada

Working papers are generally published in the language of the author, with an abstract in both official languages, and are available on the Bank's website (see bottom of page). ***Les documents de travail sont généralement publiés dans la langue utilisée par les auteurs; ils sont cependant précédés d'un résumé bilingue. On peut les consulter dans le site Web de la Banque du Canada, dont l'adresse est indiquée au bas de la page.***

2006

2006-45	The Role of Debt and Equity Finance over the Business Cycle	F. Covas and W.J. den Haan
2006-44	The Long-Term Effects of Cross-Listing, Investor Recognition, and Ownership Structure on Valuation	M.R. King and D. Segal
2006-43	Efficient Hedging and Pricing of Equity-Linked Life Insurance Contracts on Several Risky Assets	A. Melnikov and Y. Romanyuk
2006-42	Linking Real Activity and Financial Markets: The Bonds, Equity, and Money (BEAM) Model	C. Gauthier and F. Li
2006-41	An Optimized Monetary Policy Rule for ToTEM	J.-P. Cayen, A. Corbett, and P. Perrier
2006-40	Education and Self-Employment: Changes in Earnings and Wealth Inequality	Y. Terajima
2006-39	Short-Run and Long-Run Causality between Monetary Policy Variables and Stock Prices	J.-M. Dufour and D. Tessier
2006-38	Conditioning Information and Variance Bounds on Pricing Kernels with Higher-Order Moments: Theory and Evidence	F. Chabi-Yo
2006-37	Endogenous Borrowing Constraints and Consumption Volatility in a Small Open Economy	C. de Resende
2006-36	Credit in a Tiered Payments System	A. Lai, N. Chande, and S. O'Connor
2006-35	Survey of Price-Setting Behaviour of Canadian Companies	D. Amirault, C. Kwan, and G. Wilkinson
2006-34	The Macroeconomic Effects of Non-Zero Trend Inflation	R. Amano, S. Ambler, and N. Rebei
2006-33	Are Canadian Banks Efficient? A Canada–U.S. Comparison	J. Allen, W. Engert, and Y. Liu
2006-32	Governance and the IMF: Does the Fund Follow Corporate Best Practice?	E. Santor
2006-31	Assessing and Valuing the Non-Linear Structure of Hedge Fund Returns	A. Diez de los Rios and R. Garcia
2006-30	Multinationals and Exchange Rate Pass-Through	A. Lai and O. Secrieru
2006-29	The Turning Black Tide: Energy Prices and the Canadian Dollar	R. Issa, R. Lafrance, and J. Murray

Copies and a complete list of working papers are available from:
Pour obtenir des exemplaires et une liste complète des documents de travail, prière de s'adresser à :

Publications Distribution, Bank of Canada
234 Wellington Street, Ottawa, Ontario K1A 0G9
Telephone: 1 877 782-8248
(toll free in North America)
Email: publications@bankofcanada.ca
Website: http://www.bankofcanada.ca

Diffusion des publications, Banque du Canada
234, rue Wellington, Ottawa (Ontario) K1A 0G9
Téléphone : 1 877 782-8248 (sans frais en Amérique du Nord)
Adresse électronique : publications@banqueducanada.ca
Site Web : http://www.banqueducanada.ca